

## CLAIMS

We claim:

1. A method for predicting the behavior of a workload across a plurality of nodes,  
5 the method comprising:
  - a) receiving a workload to be executed;
  - b) executing the workload on a single node;
  - c) tracing the execution of the workload;
  - d) based on this execution, predicting the behavior of the workload across a  
10 plurality of nodes by identifying potential data conflicts; and
  - e) outputting the prediction.
2. The method of claim 1 wherein the action of identifying potential data conflicts  
comprises predicting how many data conflicts will occur.
3. The method of claim 1 wherein the action of identifying potential data conflicts  
15 comprise predicting types of data conflicts.
4. The method of claim 3 in which the types of data conflicts comprises a read-write  
conflict.
5. The method of claim 3 in which the types of data conflicts are based upon types  
of operations needed to resolve the data conflicts.
- 20 6. The method of claim 3 in which the different types of data conflicts have differing  
levels of expense associated with operations needed for data conflict resolution.
7. The method of claim 1 in which the potential data conflicts are at the granularity  
of a data block.

8. The method of claim 1 in which the potential data conflicts are identified based upon workload division between sessions.
9. The method of claim 1 further comprising:
- f) selecting a number of nodes;
  - g) dividing the traced execution of the workload across the number of nodes.
10. The method of claim 9 in which modulo division is used to divide the traced execution of the workload across the number of nodes.
11. The method of claim 9 in which the number of nodes corresponds to an anticipated number of nodes for a distributed computing system.
12. The method of claim 9 in which a modulo class represents a node in the number of nodes.
13. The method of claim 1 in which the potential data conflicts are used to compute costs of migrating the workload to a distributed system.
14. A method for distributing a workload across a plurality of nodes, the method comprising:
- a) receiving a workload to be executed;
  - b) executing the workload on a single node;
  - c) tracing the execution of the workload;
  - d) forming a workload distribution scheme that distributes the workload across a plurality of nodes by identifying potential data conflicts; and
  - e) outputting the workload distribution scheme.

15. The method of claim 14, wherein forming the workload distribution scheme comprises determining workload distribution in manner which reduces the potential data conflicts.

16. The method of claim 14, wherein the workload distribution scheme is based upon  
5 data accesses.

17. The method of claim 16 in which the workload is grouped in the workload distribution scheme to maximize intersection of data access on a same group of nodes.

18. The method of claim 16 in which the workload is grouped in the workload distribution scheme to minimize intersection of data access across different groups of  
10 nodes.

19. The method of claim 14, wherein the workload distribution scheme is based upon access frequencies.

20. The method of claim 19 in which data objects accessed by the workload are associated with weighting factors.

21. The method of claim 20 in which not all the data objects are associated with same  
15 weighting factors.

22. The method of claim 20 in which a weighted correlation is performed between the data objects and entities that access the data objects.

23. The method of claim 22 in which the entities that access the data objects  
20 comprises sessions.

24. The method of claim 22 in which subsets of the entities that access the data objects are grouped together.

25. The method of claim 24 in which a data structure is employed to represent an affinity between one of the entities that access the data objects and another of the entities.

26. The method of claim 14 in which the workload comprises data access upon one or more hierarchical objects.

5 27. The method of claim 26 in which tracing the execution of the workload comprises tracing identifiers for the one or more hierarchical objects.

28. The method of claim 14 in which tracing the execution of the workload comprises tracing identifiers associated with entities that access data.

29. The method of claim 28 in which the entities comprise sessions.

10 30. The method of claim 28 in which the workload distribution scheme distributes the workload based upon partitioning of the entities that access data.

31. The method of claim 30 in which an association is formed between partitioning of the entities that access data and partitioning of one or more applications within the workload.<sup>2</sup>

15 32. A computer program product that includes a medium usable by a processor, the medium comprising a sequence of instructions which, when executed by said processor, causes said processor to execute a process for optimizing the distribution of a workload across a plurality of nodes, the process comprising:

- 20
- a) receiving a workload to be executed;
  - b) executing the workload on a single node;
  - c) tracing the execution of the workload;
  - d) based on this execution, optimizing the distribution of the workload across a plurality of nodes by identifying potential data conflicts; and

- e) outputting the optimized distribution scheme.

33. A computer program product that includes a medium usable by a processor, the medium comprising a sequence of instructions which, when executed by said processor, causes said processor to execute a process for distributing a workload across a plurality of

5 nodes, the process comprising:

- a) receiving a workload to be executed;
- b) executing the workload on a single node;
- c) tracing the execution of the workload;
- d) forming a workload distribution scheme that distributes the workload

10 across a plurality of nodes by identify potential data conflicts; and

- e) outputting the workload distribution scheme.

34. A system for distributing a workload across a plurality of nodes, comprising:

- a) means for receiving a workload to be executed;
- b) means for executing the workload on a single node;
- c) means for tracing the execution of the workload;
- d) means for forming a workload distribution scheme that distributes the

workload across a plurality of nodes by identify potential data conflicts; and

- e) means for outputting the workload distribution scheme.

35. A system for optimizing the distribution of a workload across a plurality of nodes,

20 comprising:

- a) means for receiving a workload to be executed;
- b) means for executing the workload on a single node;
- c) means for tracing the execution of the workload;

- d) means for based on this execution, optimizing the distribution of the workload across a plurality of nodes by identifying potential data conflicts; and
- e) means for outputting the optimized distribution scheme.